

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions and listings of claims in the application.

1. (currently amended) A method for use in a system for storing and accessing data, the system including at least one initiator, at least one target comprising a mirrored virtual target or a physical storage device, and ~~at least one a~~ switch, the initiator, target and switch communicating using at least one network, the ~~at least one~~ switch including a plurality of ports and ~~respective processing circuitry affiliated with each respective port~~ a plurality of processing units, wherein each processing unit is associated with at least one port of said plurality of ports to provide load balancing at said at least one port, the method comprising:

providing a plurality of request paths over the network to the target from the switch, each path passing through ~~at least one port~~ one or more ports of the switch;

receiving at said ~~at least one port~~ one or more ports, a plurality of storage input/output requests; and

for each request at ~~each port~~ said one or more ports, dynamically load balancing each request amongst the paths by using ~~said respective processing circuitry affiliated with each respective port~~ one or more processing units associated with said one or more ports to determine the appropriate path.

2. (original) The method of claim 1, wherein dynamically load balancing amongst the paths includes:

determining a respective average response time for each path;

passing a request received by the switch from the initiator to the target along the path with the shortest average response time.

3. (original) The method of claim 1, wherein the target is a physical storage device.

4. (original) The method of claim 1, wherein the target is a virtual target.

5. (original) The method of claim 1, wherein the target is a mirrored target with a plurality of members and wherein load balancing amongst the paths includes:

determining a respective average response time of each member of the mirrored target;

passing a request received by the switch from the initiator to the target to the member with the shortest average response time.

6. (original) The method of claim 5, wherein the request is a read request.

7. (cancelled)

8. (currently amended) A method for use in a storage network including an initiator, a storage device, and a switch, the initiator, storage device and switch communicating via a network the switch including a plurality of ports and ~~affiliated processing circuitry with each respective port~~ a plurality of processing units, wherein each processing unit is associated with at least one port of said plurality of ports to provide load balancing at said at least one port, the method comprising:

providing a plurality of storage input/output request paths over the network from the storage device to the switch, each path passing through ~~at least one port~~ one or more ports of the switch;

receiving a plurality of storage input/output requests at ~~least one port~~ said one or more ports;

determining a respective average response time for each path using ~~said affiliated processing circuitry~~ one or more processing units associated with said one or more ports; and

passing a storage input/output request received by the switch from the initiator to the storage device along the routing path with the shortest average response time using ~~said affiliated processing circuitry~~ one or more processing units associated with said one or more ports.

9. (currently amended) A method for use in a storage network including an initiator, a mirrored virtual target having a plurality of members, and a switch, the switch including a plurality of ports ~~each with affiliated processing circuitry~~ and a plurality of processing units, wherein each

processing unit is associated with at least one port of said plurality of ports to provide load balancing at said at least one port, the method comprising:

providing a plurality of request paths from the switch to each member of the mirrored virtual target via a network, each path passing through one port of the switch;

receiving a plurality of storage input/output requests at ~~least one port~~ one or more ports of said switch;

determining a respective average response time for each path using ~~said affiliated processing circuitry~~ one or more processing units associated with said one or more ports; and

passing a storage input/output request received by the switch from the initiator to the member with the shortest average response time using ~~said affiliated processing circuitry~~ one or more processing units associated with said one or more ports.

10. (original) The method of claim 9, wherein the request is a read request.

11. (currently amended) A method for use in a storage network including a switch, a plurality of initiators, and a plurality of targets, the switch, initiators and targets communicating via at least one network, the switch including a plurality of ports each port having affiliated processing circuitry, wherein some of the targets are mirrored targets with a plurality of members and some of the targets are physical storage devices, the method comprising:

providing a plurality of storage request paths between a switch and a physical storage device from a first initiator;

providing a respective path from a second initiator to each member of a mirrored target via the switch;

determining a respective average response time for each path from the first initiator to the physical storage device and for each path from the second initiator to each member of the mirrored target;

receiving a plurality of storage input/output requests at at least one port, at least a first request being from the first initiator to the physical storage device and at least a second request from the second initiator to the member of the mirrored target;

passing the first request received by the switch from the first initiator to the physical storage device along a path to the physical storage device with the shortest average response time using said affiliated processing circuitry; and

passing the second request received by the switch from the second initiator to the member of the mirrored target along a path to the member with the shortest average response time using said affiliated processing circuitry.

12. (original) The method of claim 11, wherein:

the step of passing a first request is performed by a first linecard in the switch; and

the step of passing a second request is performed by a second linecard in the switch.

13. (original) The method of claim 11, wherein the step of passing a first request and the step of passing a second request are both performed by the same linecard.

14. (original) The method of claim 11, wherein the switch includes a plurality of linecards and wherein the step of determining is performed by each linecard.

15. (currently amended) A switch for use in a storage network, comprising:

a plurality of ports, at least one port coupled to a network and communicating with an initiator, and at least one port coupled to the network and communicating with at least one target comprising a mirrored virtual target or a physical storage device; and

load balancing circuitry associated with each port, said circuitry determining a plurality of processing units, wherein each processing unit is associated with one or more ports of said plurality of ports and determines, for storage level input/output requests received at said one or more ports, a

respective average response time for each of a plurality of storage input/output paths between the switch and the target; and

~~passing wherein the switch passes a first storage level input/output request received by the switch from the initiator at said at least one port communicating with said initiator to the target along one of said plurality of paths having a shortest average response time using a processing unit associated with said at least one port communicating with said initiator-said circuitry.~~

16. (currently amended) The switch of claim 15, wherein:

~~the load balancing circuitry each processing unit includes a storage processor and a CPU processing unit.~~

17. (currently amended) A switch for use in a storage network, the network including an initiator and a target in communication with the initiator by a plurality of paths between the switch and a target, each path passing through the switch, the switch comprising:

a plurality of ports, at least one of the plurality of paths passing through at least one of the plurality of ~~ports; ports, each port including processing circuitry associated with the port; and~~
a plurality of processing units, wherein each processing unit is associated with one or more ports of said plurality of ports to provide load balancing at said one or more ports; and

means for load balancing amongst the paths using said ~~processing circuitry~~ plurality of processing units associated with each of said ports, said means including means for receiving a plurality of storage input/output requests at said plurality of ports and passing the storage input/output request received by the switch from the initiator to the target along the path with the shortest average response time.

18. (original) The switch of claim 17, wherein the means for load balancing includes:

means for maintaining statistics for the response time of each path;

means for passing a request received by the switch from the initiator to the target along the path with the shortest average response time.

19. (currently amended) A storage network, including:
an initiator;
a target;
a switch including a plurality of ports and a plurality of processing units, wherein each processing unit is associated with at least one port to provide load balancing at said at least one port;
a plurality of storage input/output request routing paths from the switch to the target over a network, each path passing via ~~at least one port~~ one or more ports of the switch;
wherein the switch includes statistical information regarding the response time for each path;
and
wherein the switch is designed to forward a storage level input/output request from the initiator to the target along the path with the shortest response time to other storage level requests using ~~processing circuitry affiliated~~ said plurality of processing units associated with each of the ports.

20. (original) The storage network of claim 19, wherein the target is a physical storage device.

21. (original) The storage network of claim 19, wherein the target is a virtual target.

22. (original) The storage network of claim 19, wherein the target is a mirrored target with a plurality of members and wherein the plurality of paths are respective paths to each member.

23. (currently amended) ~~A machine~~ One or more computer-readable media which has instructions having computer-readable code stored thereon, which when executed by a switch in a storage network causes the switch to perform the following steps:

providing a plurality of paths to a target from the switch over a network coupling the storage switch, target and initiator, said switch including a plurality of ports and a plurality of processing

units, wherein each processing unit is associated with at least one port of said plurality of ports to provide load balancing at said at least one port, wherein each path passing of said plurality of paths passes through at least one of a plurality of ports one or more ports of the switch;

receiving at said ~~at least one port~~ one or more ports, a plurality of storage input/output requests;

determining a respective response time of each path using ~~processing circuitry affiliated with said at least one of the plurality of ports~~ one or more processing units associated with said one or more ports; and

passing the storage input/output request received by the switch from the initiator to the target along the path with the shortest average response time.

24. (original) The machine readable media of claim 23, wherein the target is a physical storage device.

25. (original) The machine readable media of claim 23, wherein the target is a virtual target.

26. (currently amended) The machine readable media of claim 23, wherein the target is a mirrored target with a plurality of members and wherein the instructions ~~[[to]]~~ further include:

determining a respective response time of each member of the mirrored target; and

passing a request received by the switch from the initiator to the target to the member with the shortest average response time.